

according to the group profile, or according to the group profile and the user's user profile when the user also has a separate user profile.

[0086] It is possible that a user may switch his interests so dramatically that his new interests and preferences bear little resemblance to his user profile, or a user may be temporarily interested in a new topic. In this case, personalized search results produced according to the embodiments depicted in **FIGS. 9A and 9B** may be less favorable than search results ranked in accordance with the generic ranks of the documents in the search results. Additionally, the search results provided to a user may not include new websites among the top listed documents because the user's profile tends to increase the weight of older websites that the user has visited (i.e., older websites from which the user has viewed or downloaded web pages) in the past.

[0087] To reduce the impact caused by a change in a user's preferences and interests, the personalized search results may be merged with the generic search results. In one embodiment, the generic search results and personalized search results are interleaved, with the odd positions (e.g., 1, 3, 5, etc.) of a search results list reserved for generic search results and the even positions (e.g., 2, 4, 6, etc.) reserved for personalized search results, or vice versa. Preferably, the items in the generic search results will not duplicate the items listed in the personalized search results, and vice versa. More generally, generic search results are intermixed or interleaved with personalized search results, so that the items in the search results presented to the user include both generic and personalized search results.

[0088] In another embodiment, the personalized ranks and generic ranks are further weighted by a user profile's confidence level. The confidence level takes into account factors such as how much information has been acquired about the user, how close the current search query matches the user's profile, how old the user profile is, etc. If only a very short history of the user is available, the user's profile may be assigned a correspondingly low confidence value. The final score of an identified document can be determined as:

$$\text{FinalScore} = \text{ProfileScore} * \text{ProfileConfidence} + \text{GenericScore} * (1 - \text{ProfileConfidence}).$$

[0089] When intermixing generic and personalized results, the fraction of personalized results may be adjusted based on the profile confidence, for example using only one personalized result when the confidence is low.

[0090] Sometimes, multiple users may share a machine, e.g., in a public library. These users may have different interests and preferences. In one embodiment, a user may explicitly login to the service so the system knows his identity. Alternatively, different users can be automatically recognized based on the items they access or other characteristics of their access patterns. For example, different users may move the mouse in different ways, type differently, and use different applications and features of those applications. Based on a corpus of events on a client and/or server, it is possible to create a model for identifying users, and for then using that identification to select an appropriate "user" profile. In such circumstances, the "user" may actually be a group of people having somewhat similar computer usage patterns, interests and the like.

#### [0091] Personalization of Advertisements

[0092] Referring again to **FIG. 1**, the content analysis module **112** receives from the personalized search results from the personalization server **108**, which then analyses the documents referenced therein, and provides a search profile to the advertisement server. The advertisement server **114** uses the search profile to select from the advertisement database **116** one or more advertisements for displaying in conjunction with the personalized search results.

[0093] The content analysis module **112** creates the search profile by determining key topic words or terms that are descriptive of the documents references in personalized search results as a group. Thus, for selected documents in the personalized search results, the content analysis module **112** determines a set of one or more topics, and then uses this set of topics to determine the topics descriptive of the personalized search results (e.g., selecting the N most frequently occurring topics, or some other filtering/selection process). The content analysis module **112** may apply any type of topic extraction methods known in the art or developed hereafter, as the particular algorithm used for topic extraction is not a limitation of the invention.

[0094] The content analysis module **112** can analyze of the documents in the personalized search results, or any subset thereof. In one embodiment, the personalized search results form a plurality of pages, each page containing some number of the documents. The documents that would be on the first page of results are the subset which the content analysis module **112** analyzes. This approach is beneficial since the documents on this first page are those most relevant to the user's interests, and hence the resulting search profile will likewise contain the most relevant terms and topics.

[0095] In one embodiment, the content analysis module **112** uses the methods described above with respect to **FIGS. 6, and 7A-7B** for constructing the term based profile of the user. Here, the operational goal is a set of terms that describe the topics of the personalized search results. In another embodiment, the content analysis module **112** uses a combination of internal document analysis that extracts topics based on the frequencies of key words in the document and in the entire document collection, and link analysis (based on the inbound and outbound link structure of each document). As a particular example of the latter, the content analysis module **112** can determine if a given document in the personalized search results is linked to one or more topics in topical directory (e.g., (<http://dmoz.org/>), and if so, uses these linked topics as candidate topics for the document. Further details of these types of methods are disclosed the Relevant Advertisements Application, cited above, which is incorporated by reference herein. In another embodiment, the content analysis module **112** uses a probabilistic model to determine the topics for inclusion in the search profile. One method of generation and use of a probabilistic model in this manner is described in the Clusters of Related Words Application, cited above, which is also incorporated by reference herein.

[0096] In any of these embodiments, the content analysis module **112** provides a search profile that includes a set of terms that describe the personalized search results, and may be characterized as the topics that the documents in the personalized search results are about. The search profile is provided to the advertisement server **114**, which then selects